

Reprint Series
28 February 1986, Volume 231, pp. 943-950

SCIENCE

Computer Networking for Scientists

DENNIS M. JENNINGS, LAWRENCE H. LANDWEBER, IRA H. FUCHS,
DAVID J. FARBER, W. RICHARDS ADRIAN

Copyright © 1986 by the American Association for the Advancement of Science

Computer Networking for Scientists

DENNIS M. JENNINGS, LAWRENCE H. LANDWEBER, IRA H. FUCHS,
DAVID J. FARBER, W. RICHARDS ADRIAN

Scientific research has always relied on communication for gathering and providing access to data; for exchanging information; for holding discussions, meetings, and seminars; for collaborating with widely dispersed researchers; and for disseminating results. The pace and complexity of modern research, especially collaborations of researchers in different institutions, has dramatically increased scientists' communications needs. Scientists now need immediate access to data and information, to colleagues and collaborators, and to advanced computing and information services. Furthermore, to be really useful, communication facilities must be integrated with the scientist's normal day-to-day working environment. Scientists depend on computing and communications tools and are handicapped without them.

A SCIENTIST SHOULD BE ABLE TO USE COMPUTING AND communications tools by working at an advanced graphics workstation. Through that single window, the scientist may gain access to required computing facilities and databases and communicate with peers, colleagues, and scholars throughout the world. This combination of computing and communications is called computer networking. Computer networks provide the base that combines geographically dispersed researchers, computing resources, and information into a single integrated computer and communications environment. Unfortunately, the development of computer networks has been fragmented and incomplete. The result has been a bewildering array of different technologies and of

D. M. Jennings is program director for networking at the National Science Foundation's Office of Advanced Scientific Computing, Washington, DC 20550. He is on leave of absence from University College, Dublin, where he is director of the University Computing Service. L. H. Landweber is professor of computer science at the University of Wisconsin, Madison 53706. I. H. Fuchs is vice president for computing and information technology at Princeton University, Princeton, NJ 08544. D. J. Farber is professor of electrical engineering and of computer science at the University of Delaware, Newark 19716. W. R. Adrián is deputy director of the Division of Computer Research at the National Science Foundation.

different and incompatible networks. The scientist has been burdened with multiple access procedures, applications software interfaces, operating systems, and data formats. However, recent developments, including the National Science Foundation's new networking program NSFnet, the emerging convergence of the community-based computer networks, and the growing focus on the adoption of standard computer networking protocols should reduce this burden. Nevertheless, the promise of the convergence of computing and communications (1)—of computer networking—remains to be fulfilled.

NSFnet

NSFnet will probably have the most impact on science of all networking activities in the United States at this time. Being based on new and existing networks, it will provide both high-speed access to supercomputers and communication between scientists in all disciplines throughout the nation. Although initially designed for supercomputer users to gain access to supercomputers and to communicate with each other, NSFnet is expected to be a general-purpose computer communications network for the whole academic research community and associated industrial researchers.

The development of NSFnet is part of the NSF supercomputer initiative. This program resulted from the growing concern in the research community over the last few years that academic research has been severely constrained by the lack of access to advanced computing facilities. Several reports (2-4) highlighted the problems: (i) large computers have become an important means of making new discoveries, (ii) there is an immediate need to make supercomputers available to U.S. researchers, and (iii) computer networks are required to link researchers to supercomputers and to each other.

In response to these concerns, NSF established the Office of Advanced Scientific Computing (OASC), which immediately initiated two programs: the supercomputer centers program to provide supercomputer cycles, and the networking program to build a national supercomputer access network—NSFnet.

In 1984-85, OASC purchased supercomputer cycles from three

Table 1. NSF supercomputer centers.

Center	Supercomputers		
	1984-85	1985-86	1987
<i>Resource</i>			
Purdue	Cyber 205	Cyber 205	
Minnesota	Cray 1A	Cray 2/ Cyber 205	
Boeing	Cray 1S	Cray X-MP/24	
AT&T Bell Labs		Cray X-MP/12	
Colorado State		Cyber 205	
Digital Productions		Cray X-MP/22	
<i>National</i>			
JVNC-Princeton	Cyber 205	ETA-10	
SDSC-San Diego	Cray X-MP/48	Cray X-MP/48	
NCSA-Illinois	Cray X-MP/24	Cray X-MP/48	
Theory-Cornell	IBM 3084/ FPS 264s	IBM 3084/ FPS	
Pittsburgh	Cray 1S	Cray 1S	

existing supercomputer centers at Purdue University, the University of Minnesota, Boeing Computer Services, AT&T Bell Laboratories, Colorado State University, and Digital Productions (Table 1). By the end of 1985, a total of 30,000 hours of supercomputer time had been allocated under this program to approximately 800 users, and more than 9000 hours had been consumed. In 1984 also, the OASC issued a project solicitation for national supercomputer centers. As a

result, four new NSF centers were funded in 1985—the John von Neumann Center (JVNC) at Princeton University, the San Diego Supercomputer Center (SDSC) on the campus of the University of California at San Diego, the National Center for Supercomputer Applications (NCSA) at the University of Illinois, and the Theory Center, a production and experimental supercomputer center at Cornell University. More recently a fifth center has been established in Pittsburgh, to be run by Westinghouse, Carnegie-Mellon University, and University of Pittsburgh (Table 1).

The NSFnet networking activities were initiated in December 1984 when a panel of the OASC confirmed that networking was a fundamental component of the supercomputer initiative, and, moreover, that a network could be designed to meet the requirements of this initiative while providing the basis for a future, general purpose, national academic research network (5). The report proposed a two-phased approach for the development of the network: phase 1 to connect supercomputer users to the supercomputer centers and to each other; and phase 2 to provide a general high-speed network, with speeds of 1.544 megabits per second (Mbps), commonly called "T1 speed," or greater. In addition, a variety of experiments to understand better how to utilize and integrate a number of network topologies and usage modalities are to be initiated.

The general strategy recommended by the networking panel report was that the NSFnet should begin by taking advantage of the existing academic networks. NSFnet should be built as a "network of networks" rather than as a separate new computer network. This general approach is based on the experience gained by the Depart-

Table 2. NSFnet. List of planned member institutions. Key: ARPANET, an existing or planned ARPANET site; SDSC, a San Diego consortium network site; JVNC, a Princeton (JVNC) consortium network site; NCAR, a National Center for Atmospheric Research (NCAR) satellite network site; Illinois, a direct 56-kbps connection to the Illinois Supercomputer Center; backbone, a supercomputer center on the NSFnet backbone network.

Institution	Network	Institution	Network
Agouron Institute	SDSC	University of North Carolina	ARPANET
University of Arizona	JVNC	North Carolina State University	ARPANET
AT&T Bell Labs, New Jersey	ARPANET	Northwestern University	ARPANET, Illinois
University of California, Berkeley	ARPANET, SDSC	Ohio State University	ARPANET
Boeing Computer Services	ARPANET	Oregon State University	NCAR
Brown University	JVNC	University of Pennsylvania	ARPANET, JVNC
California Institute of Technology	ARPANET, SDSC	Pennsylvania State University	JVNC
Carnegie-Mellon University	ARPANET	University of Pittsburgh	ARPANET
University of Chicago	Illinois	Princeton University	JVNC
Colorado State University	ARPANET, NCAR	Purdue University	ARPANET
University of Colorado	JVNC, NCAR	Rice University	ARPANET
Columbia University	ARPANET, JVNC	University of Rochester	ARPANET, JVNC
Cornell University	ARPANET, backbone	Rutgers University	ARPANET, JVNC
City University of New York	ARPANET	Salk Institute	SDSC
University of Delaware	ARPANET	San Diego Supercomputer Center	ARPANET, SDSC, backbone
Duke University	ARPANET	University of California, San Diego	SDSC
Harvard University	ARPANET, JVNC	San Diego State University	SDSC
University of Hawaii	SDSC	University of California, San Francisco	SDSC
Institute for Advanced Studies, at Princeton University	JVNC	University of California, Santa Barbara	ARPANET
University of Illinois, Urbana	ARPANET, NCAR, backbone, Illinois	Scripps Clinic and Research Foundation	SDSC
University of Illinois, Chicago	Illinois	Scripps Institute of Oceanography	SDSC
Indiana University	ARPANET, Illinois	Southwest Fisheries	SDSC
John von Neumann Center	ARPANET, JVNC, backbone	Stanford University	ARPANET, SDSC
Kitt Peak Observatory	SDSC	State University of New York, Stony Brook	ARPANET
Lawrence Berkeley Laboratory	ARPANET	University of California, Los Angeles	ARPANET, SDSC
University of Maryland	ARPANET, SDSC, NCAR	University of Texas, Austin	ARPANET
University of Miami	NCAR	University of Utah	ARPANET, SDSC
University of Michigan	ARPANET, SDSC, NCAR	University of Washington	ARPANET, SDSC, NCAR
University of Minnesota	ARPANET	Westinghouse (Pittsburgh)	ARPANET, backbone
Massachusetts Institute of Technology	ARPANET, JVNC	University of Wisconsin	ARPANET, SDSC, NCAR
National Center for Atmospheric Research	ARPANET, NCAR, backbone	Woods Hole Oceanographic Institution	NCAR
National Science Foundation	ARPANET	Yale University	ARPANET
New York University	JVNC		

ment of Defense Advanced Research Projects Agency (DARPA) when developing the ARPANET network. The ARPANET experience demonstrated that an Internet, a collection of networks with the same higher level protocols, could provide access to remote computing resources from within a researcher's own local computing environment.

A common set of networking protocol standards has, therefore, to be adopted by NSF in order to build the NSFnet Internet. NSF has decided on the ARPANET protocol (TCP-IP and associated application protocol—the DARPA protocol suite) as the initial NSFnet standards. A migration to the emerging international standards organization (ISO) open systems interconnection (OSI) networking protocol standards (6) is to be undertaken as these become available.

In the initial (current) phase of the development of NSFnet the network will be based on various component networks. These include the wide-area community networks (in particular ARPANET), the supercomputer center consortia networks (at JVNC and SDSC), a supercomputer "backbone" network, the various state networks, and, most important, the campus networks. In addition a number of pilot networking projects will be undertaken.

The goal of phase 1 is to provide supercomputer access and to support communication between supercomputer researchers. By September 1986, on the basis of the developments planned to date and of the expansion of the ARPANET, more than 60 major research universities in the United States are expected to be connected to NSFnet, to the NSF supercomputers, and to each other (Table 2). Other institutions are expected to be added to this list during 1986.

The average researcher working at a terminal or workstation at one of these institutions will then be able to connect to and use various computer resources—including the NSF supercomputer centers—to run interactive and batch jobs, receive output, transfer files, and use the electronic mail facilities to communicate with any colleague throughout the nation. Typically, an individual researcher will have either a terminal connected to a local super-minicomputer or a graphics workstation. These computers will be connected to a local area network (LAN), that will provide local communications and resource sharing. It is expected that this LAN will itself be connected to other LAN's on the campus, and that the collection of interconnected LAN's will form a campus network—with, ideally, a campus-wide service organization taking responsibility for the overall network services provided. In turn, this campus network will be connected, via a campus gateway system, to one or more of the wide-area networks in the NSFnet to provide the researcher with computer communications across the United States.

Wide-Area Networks

In the early 1970's the United States took the lead in the development of wide-area networks for academic research. A variety of networks and network technologies have been developed. In this section, some of these are outlined.

ARPANET. The ARPANET, which is a major component of the NSFnet, began in 1969 as an R&D project managed by DARPA. ARPANET was an experiment in resource sharing, and provided survivable (multiply connected), high bandwidth [56 kilobits per second (kbps)] communications links between major existing computational resources and computer users in academic, industrial, and government research laboratories (7) (Fig. 1). ARPANET is managed and funded by the Defense Communications Agency (DCA) with user services provided by a network information center at SRI International.

ARPANET served as a test for the development of advanced

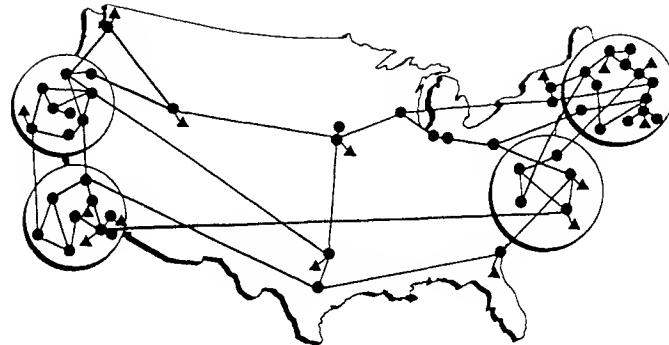


Fig. 1. The 1985 Configuration of the Defense Advanced Research Projects Agency (DARPA) network, ARPANET. The ARPANET is both an advanced networking R&D testbed for DARPA and an operational network supporting many DARPA-sponsored researchers in universities, national laboratories, and industry. In October 1985, NSF reached agreement with DARPA to expand the ARPANET by approximately 40 sites for use by NSF-sponsored supercomputer users. The network is built of 56-kbps links between (●) interface message processors (IMPs), to which host computers are connected, and (▲) terminal access controllers (TAC's). Services provided include remote terminal access, file transfer, and electronic mail. Major clusters of host connections are in Boston, Washington, DC, San Francisco, and Los Angeles. [Courtesy of DARPA]

network protocols including the TCP-IP protocol suite introduced in 1981. TCP-IP and particularly IP, the internet protocol (8, 9), introduced the idea of internetworking—allowing networks of different technologies and connection protocols to be linked together while providing a unified internetwork addressing scheme and a common set of transport and application protocols. This development allowed networks of computers and workstations to be connected to the ARPANET, rather than just single-host computers. TCP-IP remain the most available of the advanced, non-vendor-specific, networking protocols and have strongly influenced the current international standards activity. TCP-IP provide a variety of application services, including remote logon (Telnet), file transfer (FTP), and electronic mail (SMTP and RFC822).

ARPANET technology was so successful that in 1982, the Department of Defense (DOD) abandoned their AUTODIN II network project and adopted ARPANET technology for the Defense Data Network (DDN). The current MILNET, which was split from the original ARPANET in 1983, is the operational, unclassified network component of the DDN, while ARPANET remains an advanced network R&D testbed for DARPA. In practice, ARPANET has also been an operational network supporting DOD, the Department of Energy (DOE), and some NSF-sponsored computer science researchers. This community has come to depend on the availability of the network. Until the advent of NSFnet, access to ARPANET was restricted to this community.

As an operational network in the scientific and engineering research community, and with the increasing availability of affordable super-minicomputers, ARPANET was used less as a tool for sharing remote computational resources than it was for sharing information. The major lesson from the ARPANET experience is that information sharing is a key benefit of computer networking. Indeed it may be argued that many major advances in computer systems and artificial intelligence are the direct result of the enhanced collaboration made possible by ARPANET.

However, ARPANET also had the negative effect of creating a have-have not situation in experimental computer research. Scientists and engineers carrying out such research at institutions other than the twenty or so ARPANET sites were at a clear disadvantage



Fig. 2. The 1985 Configuration of the Computer Science Research Network, CSNET; which has three major components: (■) ARPANET sites; (◆) X25NET sites connected to the public X.25 data networks Telenet and UNINET; (●) Phonenet sites with dial-up connections to a central mail relay service at the CSNET Coordination and Information Center (CIC) run by Bolt, Beranek, and Newman (BBN). CSNET provides remote terminal access, file transfer, and electronic mail services to ARPANET and X25NET sites. Electronic mail is the only service available to Phonenet sites. [Courtesy of the CSNET CIC]

in accessing pertinent technical information and in attracting faculty and students.

In October 1985, NSF and DARPA, with DOD support, signed a memorandum of agreement to expand the ARPANET to allow NSF supercomputer users to use ARPANET to access the NSF supercomputer centers and to communicate with each other. The immediate effect of this agreement was to allow all NSF supercomputer users on campuses with an existing ARPANET connection to use ARPANET. In addition, the NSF supercomputer resource centers at Purdue University and the University of Minnesota, and the national centers at the University of Illinois and Cornell University are connected to ARPANET. In general, the existing ARPANET connections are in departments of computer science or electrical engineering and are not readily accessible by other researchers. However, DARPA has requested that the campus ARPANET coordinators facilitate access by relevant NSF researchers (Table 2).

As part of the NSFnet initiative, a number of universities have requested connection to ARPANET. Each of these campuses has undertaken to establish a campus network gateway accessible to all campus researchers, thus ensuring that individual researchers will, in due course, be able to use the ARPANET to access the NSF supercomputer centers, from within their own local computing environment (Table 2). Additional requests for connection to the ARPANET are being considered by NSF.

CSNET. Establishment of a network for computer science research was first suggested in 1974, by the NSF advisory committee for computer science. The objective of the network would be to support collaboration among researchers, provide resource sharing, and, in particular, support isolated researchers in the smaller universities.

In the spring of 1980, CSNET, the computer science network, was defined and proposed to NSF as a logical network made up of several physical networks (10) of various power, performance, and cost. NSF responded with a 5-year contract for development of the network under the condition that CSNET was to be financially self-supporting by 1986. Initially CSNET was a network with five major components—ARPANET, Phonenet (a telephone-based message-relaying service) (11), X25Net (support for the TCP-IP protocol

suite over X.25-based public data networks), a public host (a centralized mail service), and a name server (an on-line database of CSNET users to support transparent mail services). The common service provided across all these networks is electronic mail, which is integrated at a special service host, which acts as an electronic mail relay between the component networks. Thus CSNET users can send electronic mail to all ARPANET users and vice versa. CSNET, with DARPA support, installed ARPANET connections at the CSNET development sites at the universities of Delaware and Wisconsin and Purdue University.

In 1981, Bolt, Beranek, and Newman (BBN) contracted to provide technical and user services and to operate the CSNET Coordination and Information Center. In 1983, general management of CSNET was assumed by UCAR—the University Corporation for Atmospheric Research, with a subcontract to BBN. Since then, CSNET has grown rapidly and is currently an independent, financially stable, and professionally managed service to the computer research community (Fig. 2). In the beginning, the need for CSNET was not universally accepted within the computer science community. However, the momentum created by CSNET's initial success caused the broad community support it now enjoys. More than 165 university, industrial, and government computer research groups now belong to CSNET (12).

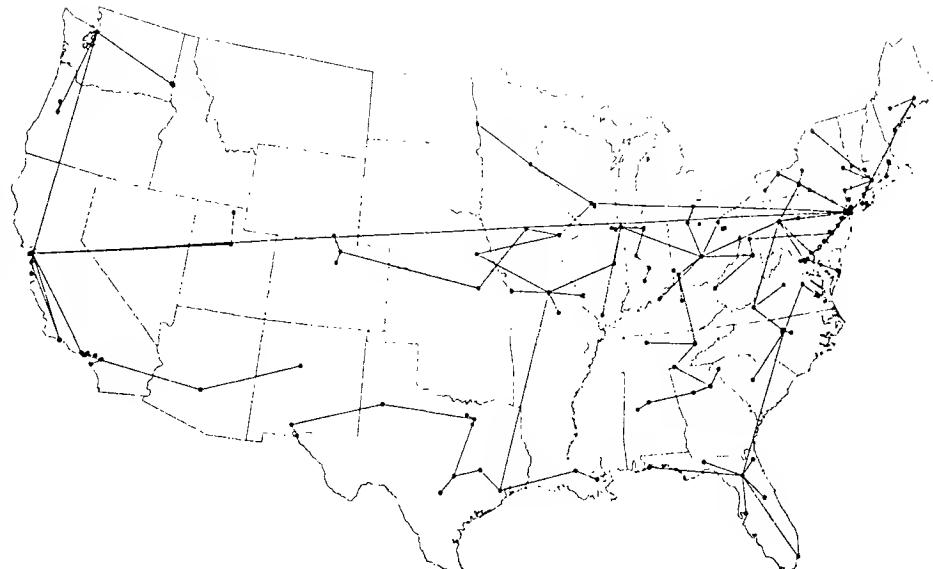
A number of lessons may be learned from the CSNET experience (12). (i) The network is now financially self-sufficient, showing that a research community is willing to pay for the benefits of a networking service. (Users pay usage charges plus membership fees ranging from \$2,000 for small computer science departments to \$30,000 for the larger industrial members.) (ii) While considerable benefits are available to researchers from simple electronic mail and mailing list services—the Phonenet service—most researchers want the much higher level of performance and service provided by the ARPANET. (iii) Providing a customer support and information service is crucial to the success of a network, even (or perhaps especially) when the users are themselves sophisticated computer science professionals. Lessons from the CSNET experience will provide valuable input to the design, implementation, provision of user services, and operation and management of NSFnet, and, in particular, to the development of the appropriate funding model for NSFnet.

CSNET, with support from the NSFnet program, is now developing the CYPRESS project which is examining ways in which the level of CSNET service may be improved, at low cost, to research departments. CYPRESS will use the DARPA protocol suite and provide ARPANET-like service on low-speed 9600-bit-per-second (bps) leased line telephone links. The network will use a nearest neighbor topology, modeled on BITNET, while providing a higher level of service to users and a higher level of interoperability with the ARPANET. The CYPRESS project is designed to replace or supplement CSNET use of X.25 public networks, which has proved excessively expensive. This approach may also be used to provide a low-cost connection to NSFnet for smaller campuses.

BITNET. In 1981, City University of New York (CUNY) surveyed universities on the East Coast of the United States and Canada, inquiring whether there was interest in creating an easy-to-use, economical network for interuniversity communications. The response was positive. Many shared the CUNY belief in the importance of computer-assisted communication between scholars. The first link of the new network, called BITNET, was established between CUNY and Yale University in May 1981.

The network technology chosen for BITNET was determined by the availability of the RSCS software on the IBM computers at the initial sites. [The name BITNET stands for Because It's Time NETwork (13).] The RSCS software is simple but effective, and

Fig. 3. The 1985 BITNET configuration. BITNET is a store-and-forward network with files and messages sent from host computer to host computer across the network. Services provided include electronic mail, file transfer, and remote job entry. The standard BITNET links are leased telephone lines running at 9600 bps. Electronic mail relays at the University of California at Berkeley and at the University of Wisconsin-Madison provide communication between BITNET, ARPANET, and CSNET users. [Courtesy of Texas A&M University]



most IBM VM-CMS computer systems have it installed for local communications, supporting file transfer and remote job entry services. The standard BITNET links are leased telephone lines running at 9600 bps. Although all the initial nodes were IBM machines in university computer centers, the network is in no way restricted to such systems. Any computer with an RSCS emulator can be connected to BITNET. Emulators are available for Digital Equipment Corporation (DEC) VAX-VMS computer systems, for VAX-UNIX systems, and for Control Data Corporation Cyber systems and others. Today, more than one-third of the computers on BITNET are non-IBM systems.

BITNET is a store-and-forward network with files and messages sent from computer to computer across the network. It provides electronic mail, remote job entry, and file transfer services, and supports an interactive message facility and a limited remote logon facility. Most BITNET sites use the same electronic mail procedures and standards as the ARPANET, and as a result of the installation of electronic mail gateway systems at the University of California at Berkeley and at the University of Wisconsin-Madison, most BITNET users can communicate electronically with users on CSNET and the ARPANET.

BITNET has expanded extremely rapidly—a clear indication that it is providing service that people need and want. The simplicity of connection to the network—acquiring a 9600-bps leased line to the nearest neighboring computer node and installing an additional line interface and a modem—provides the service at the right price. By the end of 1985 the number of computers connected was expected to exceed 600, at more than 175 institutions of higher education throughout the United States (Fig. 3). BITNET is open without restriction to any college or university. It is not limited to specific academic disciplines, and may be used for any academic or administrative purpose. However, use for commercial purposes is prohibited. In special cases, connection of commercial organizations may be sponsored by universities. A particular case is the connection of Boeing Computer Services to BITNET, as part of the NSFnet initiative, to provide remote job entry services to their Cray X-MP/24 to NSF supercomputer grantees who have access to BITNET.

Until recently BITNET had no central management structure, and was coordinated by an executive board consisting of members from the major institutions participating. This worked because most of the computers connected were managed and operated by professional service organizations in university computer centers. Howev-

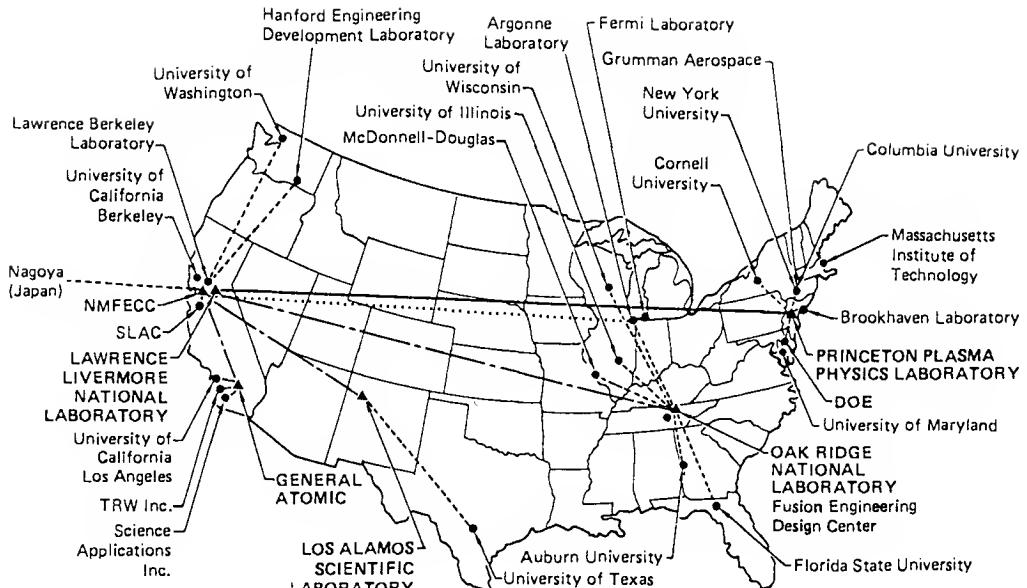
er, the growth in the network made it impossible to continue in this ad hoc fashion, and a central support organization was established with support from an IBM grant. The central support organization, called the BITNET network support center (BITNSC), has two parts: A user services organization, the network information center (BITNIC), which provides user support, a name server and a variety of databases, and the development and operations center (BITDOC) to develop and operate the network. A major question facing the members of BITNET is how the funding of this central organization will be continued when the IBM grant expires in 1987.

BITNET, with support from the NSFnet Program, is now examining ways to provide ARPANET-like services to existing BITNET sites. The project, which is similar to the CSNET CYPRESS project, will explore a strategy to provide an optional path to the use of the TCP-IP procedures on existing 9.6-kbps leased lines. The possibility of upgrading these lines to multiple alternate links, providing higher reliability and availability, or to higher speed 56-kbps links is also being studied. The project will offer a higher level of service to BITNET sites choosing this path and also enable a low-cost connection to NSFnet.

MFENET. The DOE's magnetic fusion energy research network (MFENET) was established in the mid-1970's to support access to the MFE Cray 1 supercomputer at the Lawrence Livermore National Laboratory. The network uses 56-kbps satellite links, and is designed to provide terminal access to the Cray time-sharing system (CTSS), also developed at the Livermore Laboratory. The network currently supports access to Cray 1, Cray X-MP/2, Cray 2, and Cyber 205 supercomputers. The network uses special-purpose networking software developed at Livermore, and, in addition to terminal access, provides file transfer, remote output queuing, and electronic mail, and includes some specialized application procedures supporting interactive graphics terminals and local personal computer (PC)-based editing. Access to the network is in general restricted to DOE-funded researchers. Recently the network has been expanded to include the DOE-funded supercomputer at Florida State University. MFENET (Fig. 4) is funded by DOE and managed by Livermore.

MFENET has been successful in supporting DOE supercomputer users. However, the specialized nature of the communications protocols is now creating difficulties for researchers who need advanced graphics workstations that use the UNIX BSD 4.2 operating system and the TCP-IP protocols on LAN's. For these

Fig. 4. The 1985 Configuration of DOE's Magnetic Fusion Energy researchers network (MFENET). The network uses dual satellite links at 112 kbps (solid line) and 56 kbps (dashed lines) and terrestrial links at 56 kbps (dotted lines) and 9600 bps (short dashes). The network was developed at the Lawrence Livermore National Laboratory to provide access to supercomputers running the CTSS, also developed at the Livermore Laboratory. The network uses special-purpose networking software developed at MFE. Services include terminal access, file transfer, remote output queuing, and electronic mail. Abbreviations: SLAC, Stanford Linear Accelerator site; NMFECC, National Magnetic Fusion Energy Computer Center. [Courtesy of NMFECC]



and other reasons, DOE is examining how best to migrate MFENET to the TCP-IP, and later to the OSI, protocols.

The combination of the CTSS operating system and the MFENET protocols creates an effective interactive computing environment for researchers using Cray supercomputers. For this reason, two of the new NSF national supercomputer centers—San Diego (SDSC) and Illinois—have chosen the CTSS operating system. In SDSC's case, the MFENET protocols have also been chosen to support the SDSC Consortium network. In Illinois's case, a project to implement the TCP-IP protocols for the CTSS operating system has been funded by the NSFnet program, and these developments will be shared with SDSC (and with DOE) to provide a migration path for the SDSC Consortium network.

UUCP and USENET. The UUCP network was started in the 1970's to provide electronic mail and file transfer between UNIX systems. The network is a host-based store and forward network using dial-up telephone circuits and operates by having each member site dial-up the next UUCP host computer and send and receive files and electronic mail messages. The network uses addresses based on the physical path established by this sequence of dial-up connections. UUCP is open to any UNIX system which chooses to participate. There are "informal" electronic mail gateways between UUCP and ARPANET, BITNET, or CSNET, so that users of any of these networks can exchange electronic mail.

USENET is a UNIX news facility based on the UUCP network that provides a news bulletin board service. Neither UUCP nor USENET has a central management; volunteers maintain and distribute the routing tables for the network. Each member site pays its own costs and agrees to carry traffic. Despite this reliance on mutual cooperation and anarchic management style, the network operates and provides a useful, if somewhat unreliable, and low-cost service to its members. Over the years the network has grown into a worldwide network with thousands of computers participating.

Other Wide-Area Networks

Of necessity this discussion of wide-area networks has been incomplete: Other networks of interest include the Space Plasma Analysis Network (SPAN)—a network of DEC VAX computers using 9.6-kbps links and the DECNET protocols for National

Aeronautics and Space Administration's researchers; the planned Numerical and Atmospheric Sciences (NAS) network centered at Ames Research Center—a network that is expected to use existing and planned NASA communications links and the TCP-IP protocols; and the planned high-energy physics network—a network based largely on VAX computers and using the standard X.25 network level protocols plus the so-called "coloured books" protocols developed in the United Kingdom. Also, many high-energy physicists, at the Stanford Linear Accelerator, at the Lawrence Berkeley Laboratory, and at Fermi Laboratory, among others, have used DECNET to connect their DEC VAX computers together.

State Networks

A number of states have over the years developed state-wide networks to provide access to shared computing facilities and to support exchange of information among researchers. The best known of these is the Merit Computer Network in Michigan, which links the campuses of the University of Michigan and of Oakland, Michigan State, Wayne State, and West Michigan universities. This is an extensive network, providing terminal access to a wide variety of resources, and is based on the use of the X.25 network level protocols.

Other states are beginning to examine the development of a state-wide research network. An example is the proposal for a New York State education and research network (NYSERNet). This network is envisaged by the proposers to provide a computer communications infrastructure for both the academic research institutions, and for high technology industrial research laboratories in the state. The network is designed not only to support the development of research activities between the academic researchers and existing industry, but also to provide the basis for the attraction of new high-technology industry to the state.

NYSERNet is to be based on multiple redundant T1 (1.544 Mbps) links, and high-performance switches, with gateways to every campus. The network will support the DARPA protocol suite, and the host and campus gateways will run the TCP-IP protocols. The plan envisions that each campus will install a campus-wide network—a model that is entirely consistent with the NSFnet model—and that each individual researcher will be equipped with a powerful

graphics workstation. All computing and information resources on the network, including the new NSF national supercomputer center at Cornell, will be accessible from those workstations. NYSERNet, will also be gatewayed to the NSFnet, and will become an integral part of the evolving national research network.

Supercomputer consortia and "Backbone" networks, NSFnet pilot projects. Two of the NSF national supercomputer centers are consortia endeavors. The JVNC center was proposed by the Princeton consortium, and the SDSC center by the San Diego consortium. Each proposed a network to link the members of their consortium to their supercomputer center.

The Princeton consortium network. The Princeton consortium comprises 13 schools, mostly along the East Coast of the United States (Table 2). The planned consortium network is a star network linking the member campuses to the JVNC. The network uses T1 circuits (1.544 Mbps) in most cases, and each link will be terminated at a campus gateway system, providing connection to a campus-wide network—a model consistent with the NSFnet model. The campus gateway systems and the front-end computers at the JVNC will run the DARPA protocol suite, so the Princeton consortium network is, in fact, an integral part of the NSFnet. Researchers on the consortium campuses will be able to access the JVNC Cyber 205 (and by mid-1987 the ETA-10 system), and, via the consortium network, the other national supercomputer centers and the other campuses on NSFnet, from within their own local computing environments. The Princeton consortium network should be operational by June 1986.

The San Diego consortium network. The San Diego consortium comprises 19 schools, mostly along the West Coast of the United States (Table 2). The consortium network is also a star network linking the consortium member campuses to the San Diego center. The network uses 56-kbps circuits, of various types, with each link terminated at a campus remote user access system (RUAC), providing access to the supercomputer for campus researchers—a model somewhat similar to the NSFnet model. Because the SDSC will operate a CRAY X-MP/48 system running the CTSS operating system, the consortium network will initially use the MFENET protocols providing terminal access, file transfer, remote output queuing, interactive graphics, and electronic mail. Although the network will not be an integral part of the NSFnet, a migration to the DARPA protocol suite is planned and is expected to take place during 1987. As an interim measure a gateway/relay system will be installed at the SDSC, which will be accessible to the consortium users, and which will be connected to the NSFnet. Thus consortium users will be able to access the other national supercomputer centers, and other users on the NSFnet will be able to access the SDSC. The San Diego consortium network should be completed by August 1986.

The supercomputer "backbone" network. To connect the supercomputer consortia networks to all the NSF national supercomputer centers, including the long-established National Center for Atmospheric Research (NCAR) and to facilitate cooperation between the centers (such as for file transfer, data sharing, or load balancing), NSF is installing a supercomputer "backbone" network, as part of the development of NSFnet (Fig. 5). Initially, this network will be based on multiple 56-kbps circuits, with low-speed switches and gateways, but it is envisioned that the network will be upgraded to T1 circuits as the volume of user to supercomputer traffic and file-transfer traffic between supercomputer centers grows. This backbone will be integral to the NSFnet internet. The network may be expanded to include connections to other supercomputer centers and to the larger campuses.

NSFnet pilot projects. In addition to the CSNET CYPRESS project, the BITNET migration project, and the Illinois project to

develop the TCP-IP procedures for the CTSS operating system, the NSFnet program will include a number of pilot networking projects. The objective will be to explore the use of new networking technologies and to gain experience to assist with the design of the phase 2 NSFnet.

Although it is expected that several substantial projects will be funded over the next few years, to date only one pilot project has been funded—the NCAR satellite experiment. This project will utilize Ku-band (12 to 14 GHz) satellite equipment developed by the Vitalink Corporation to link together Ethernets in several locations in the United States. The central or "hub" site will be located at NCAR in Boulder, Colorado, and will broadcast at 224 kbps to several remote sites [the universities of Illinois, Maryland, Miami, Michigan, and Wisconsin, Oregon State University, and the Woods Hole Oceanographic Institution (Table 2)]. Each remote site will be able to receive data addressed to it by the hub site at up to 224 kbps, and each will have a dedicated 56-kbps return satellite path to NCAR. In addition, 56-kbps terrestrial links will be installed to Colorado University and Colorado State University. The Ku-band Earth stations used are relatively inexpensive.

The objective of the NCAR pilot project is to explore the use of the shared broadcast channel to provide high-speed communications to remote supercomputer users, to investigate the optimization required to efficiently use the satellite network with the TCP-IP protocols, and to develop the experience necessary to evaluate the more extensive use of Ku-band satellite channels and the Vitalink technology in the phase 2 NSFnet.

Campus Networks

The same factors that have motivated the development of wide-area networks—access to a variety of computing facilities and communication amongst researchers—have also motivated the development of campus networks. Until recently, these developments



Fig. 5. Planned backbone network connecting NSF-sponsored supercomputers at Cornell University, the John von Neumann Center, at the University of Pittsburgh, the University of Illinois, the National Center for Atmospheric Research, and the San Diego Supercomputer Center. The links will be 56-kbps terrestrial digital circuits connecting network gateways at each site. The supercomputer front-end computers will run the NSFnet standard protocols (TCP-IP and associated application protocols). The NSFnet backbone network will be connected to the ARPANET, to various regional and state networks, and to the planned NSF supercomputer center networks to provide NSF-sponsored supercomputer users with access to all the NSF supercomputer centers. [Courtesy of the NSF's Office of Advanced Scientific Computing]

were fragmented and uncoordinated. However, many universities have now realized that a unified approach to networking is required to provide campus-wide computer communications, improved access to central campus computer resources and campus-wide access to the wide-area networks. The details of how campus networks are being developed vary widely, but, in general, campus networks have three components: a traditional terminal network to central time-sharing systems, a variety of departmental LAN's, and a campus backbone network.

The traditional terminal network is usually based on twisted-pair telephone wiring supporting direct terminal access connections to mainframe or super-minicomputer time-sharing systems, at speeds of between 1200 bps and 9.6 kbps. Where multiple host computers are supported, some form of terminal switch-contention unit or terminal concentrators have been used to provide user selection of, and contention for, the requested service. Where terminals have been replaced by PC's and workstations, access to the time-sharing systems is provided by terminal emulation software. Transfer of files between the PC's and the central systems is provided by file-transfer programs such as the "Kermit" program, developed initially at Columbia University.

In addition to using central computer systems, many departments installed their own local computers (mini-computers, super-minis, PC's, and workstations) and purchased LAN's to connect these and to integrate departmental computing facilities. These departmental LAN's may be based on a variety of technologies (such as contention bus or token ring) and typically operate at speeds of 10 Mbps. They provide high-speed connections between the various computers, supporting electronic mail, file transfer and remote logon, plus sharing of various types of resources (files, printers, computer cycles, electronic mail, mailing lists, news, and so forth). A variety of communications procedures may be used, but the most popular are the TCP-IP protocols, because of their availability on super-minicomputers, advanced workstations, and IBM PC's.

The approach generally taken to building a campus network is to take advantage of the existence of the departmental LAN's and to build a "network of networks," by installing a "backbone," which interconnects the departmental LAN's and the centrally supported systems. The "backbone" may itself be a network or a backbone medium (such as cable television or fiber optic cable) used to physically link the networks together. Although attention tends to be focused on the physical media used, the really important issues are at the higher level communications procedures.

Because many of the computers installed on departmental LAN's already use the TCP-IP protocol suite, many universities have adopted these protocols as the campus standard. By installing IP gateways between the departmental LAN's and the backbone network, they are building campus-internets, based on the DARPA internet model. The functions provided across the campus network are an extension of those provided on departmental LAN's—high-speed connections between computers, supporting file transfer and remote logon, and a variety of servers. Where several incompatible networking protocols (DECNET, TCP-IP, Xerox XNS protocols, and so forth) have been used on departmental LAN's, application gateways or relays may be installed to interconnect the LAN's. The latter method of interconnection cannot, in general, provide the same level of user functionality as is possible with a single set of protocols. Hence, most campuses attempt to establish a single set of protocols, to achieve the maximum functionality possible across the campus. To date the majority have chosen to use the TCP-IP protocols.

The development of the campus network, and the connection of these networks to the wide-area networks, provide researchers with access to computer and information resources across the United

States. Thus the campus network is a basic component of NSFnet and of any future national academic research network.

A national academic research network. For NSF supercomputer users, the development of the phase 1 NSFnet will open the possibility of integrating computing and communications and will greatly enhance their research environment. Further enhancements will be achieved when the phase-2 NSFnet provides additional connectivity, increased network bandwidth and performance, and improved functionality. However, although NSFnet is a nationwide network and a major advance for NSF supercomputer users, it is but one step toward the vision of a national network for all research scientists and engineers.

A national academic research network involves more than just connectivity and access to remote resources. Our vision of this network is of a vast network of networks interconnecting the scientists local advanced graphics workstation environment to other local and national resources (14). Scientists and engineers will be able to work at such workstations, using tools that are both comfortable and familiar and interacting with an environment that reflects a model of their scientific world. Through the network, a researcher will be able to build programs, execute and modify models, and collect and analyze data without concern for where the tools, programs, or models reside. The scientist will be able to bring powerful computational resources to bear on problems, without explicit knowledge of the physical machines and communications involved. The procedures and formats for accessing these resources and other services will be as uniform as possible.

This means that the National Academic Research Network has to provide both the network and the software tools and application protocols to make the scientist's workstation an integral part of the larger networked environment. Our vision is of a network integrating the computer resources available and presenting these resources to the user as a single interactive system.

The NSFnet experience has already indicated the approach to be taken to develop such a National Academic Research Network. The network will be an internet—a network of networks—and technical decisions on the adoption of common networking standards will have to be made. Building this national internet will be even more complex than in the case of NSFnet because of the many research disciplines and sponsoring agencies involved. The management and organization of the national internet by the research community and by the funding agencies will also take time to coordinate and develop.

REFERENCES AND NOTES

1. D. J. Farber and P. Baran, *Science* 195, 1166 (1977).
2. F. Press, "A report on the computational needs for physics" (National Academy of Sciences, Washington, DC, 1981).
3. P. D. Lax, "Report of the panel on large scale computing in science and engineering" (National Science Foundation, Washington, DC, 1982).
4. K. K. Curtis and M. Bardon, "A national computing environment for academic research" (National Science Foundation, Washington, DC, 1983).
5. W. R. Adrion, D. J. Farber, F. F. Kuo, L. H. Landweber, J. B. Wyatt, "A report on the evolution of a national supercomputer access network: Sciencenet" (National Science Foundation, Washington, DC, 1984).
6. "Special Issue on Open Systems Interconnection (OSI)," *Proc. IEEE* 71, no. 12 (1983).
7. J. M. McQuillan and D. C. Walden, *Comput. Networks* 1, 243 (1977).
8. V. Cerf and R. Kahn, *IEEE Trans. Commun.* 22, 637 (1974).
9. J. Reynolds and J. Postel, "Official ARPA internet protocols (RFC 944)" (Network Working Group, Information Sciences Institute, Marina del Rey, CA, 1985).
10. E. H. Crocker, E. S. Szurkowski, D. J. Farber, *Proceedings of the 6th Data Communications Symposium, IEEE* (1979), pp. 18-25.
11. L. H. Landweber and M. H. Solomon, "Use of multiple networks in CSNET," *Proceedings COMPCON Spring Feb.* 22, *IEEE*, pp. 398-402.
12. D. Comer, *Commun. ACM* 26, 747 (1983).
13. I. H. Fuchs, *Perspect. Comput.* 3, 16 (1983).
14. W. R. Adrion, D. J. Farber, L. H. Landweber, *Proc. IEEE 1st Annual Conf. Supercomputing*, St. Petersburg, FL, 16 to 20 December 1985.